

Accurate Estimators for Improving Minwise Hashing and b-Bit Minwise Hashing

Ping Li

Department of Statistical Science
Cornell University
Ithaca, NY 14853
pingli@cornell.edu

Christian König

Microsoft Research
Microsoft Cooperation
Redmond, WA 98052
chrisko@microsoft.com

Abstract

Minwise¹ hashing is the standard technique in the context of search and databases for efficiently estimating set (e.g., high-dimensional 0/1 vector) similarities. Recently, b-bit minwise hashing was proposed which significantly improves upon the original minwise hashing in practice by storing only the lowest b bits of each hashed value, as opposed to using 64 bits. b-bit hashing is particularly effective in applications which mainly concern sets of high similarities (e.g., the *resemblance* > 0.5). However, there are other important applications in which not just pairs of high similarities matter. For example, many learning algorithms require all pairwise similarities and it is expected that only a small fraction of the pairs are similar. Furthermore, many applications care more about *containment* (e.g., how much one object is contained by another object) than the *resemblance*. In this paper, we show that the estimators for minwise hashing and b-bit minwise hashing used in the current practice can be systematically improved and the improvements are most significant for set pairs of low resemblance and high containment.

1 Introduction

Computing the size of set intersections is a fundamental problem in information retrieval, databases, and machine learning. For example, binary document vectors represented using w -shingles can be viewed either as vectors of very high dimensionality or as sets. The seminal work of *minwise hashing* [2, 4] is a standard tool for efficiently computing *resemblances* (Jaccard similarity) among extremely high-dimensional (e.g., 2^{64}) binary vectors, which may be documents represented by w -shingles (w -grams, w contiguous words)

with $w = 5$ or 7 [2, 4]. Minwise hashing has been successfully applied to a very wide range of real-world problems especially in the context of search; a partial list includes [2, 4, 1, 13, 7, 27, 5, 29, 17, 10, 8, 14, 19, 25].

The *resemblance*, R , is a widely used measure of similarity between two sets. Consider two sets $S_1, S_2 \subseteq \Omega = \{0, 1, 2, \dots, D-1\}$, where D , the size of the dictionary, is often set to be $D = 2^{64}$ in industry practice. Denote $a = |S_1 \cap S_2|$. R is defined as

$$R = \frac{|S_1 \cap S_2|}{|S_1 \cup S_2|} = \frac{a}{f_1 + f_2 - a}, \quad f_1 = |S_1|, \quad f_2 = |S_2|.$$

Minwise hashing applies a random permutation $\pi : \Omega \rightarrow \Omega$ on S_1 and S_2 . Based on an elementary probability result:

$$\Pr(\min(\pi(S_1)) = \min(\pi(S_2))) = \frac{|S_1 \cap S_2|}{|S_1 \cup S_2|} = R, \quad (1)$$

one can store the smallest elements under π , i.e., $\min(\pi(S_1))$ and $\min(\pi(S_2))$, and then repeat the permutation k times to estimate R . After k minwise independent permutations, $\pi_1, \pi_2, \dots, \pi_k$, one can estimate R without bias, as:

$$\hat{R}_M = \frac{1}{k} \sum_{j=1}^k 1\{\min(\pi_j(S_1)) = \min(\pi_j(S_2))\}, \quad (2)$$

$$\text{Var}(\hat{R}_M) = \frac{1}{k} R(1 - R). \quad (3)$$

The common practice is to store each hashed value, e.g., $\min(\pi(S_1))$ and $\min(\pi(S_2))$, using 64 bits [12]. The storage cost (and consequently the computational cost) will be prohibitive in large-scale applications [24].

It is well-understood in practice that one can reliably replace a permutation with a reasonable hashing function; see the original minwise hashing paper [2] and the followup theoretical work [3]. In other words, there is no need to store these k permutations.

¹First draft in March, slightly modified in June, 2011.

In this paper, we first observe the standard practice of minwise hashing, i.e., using (2), can be substantially improved for important scenarios. In fact, we will show that (2) is optimal only when the sets are of the same size, i.e., $f_1 = f_2$, which is not too common in practice.

Figure 1 presents an example based on the *webspam* dataset (available from the LibSVM site), which contains 350000 documents represented using binary vectors of $D = 16$ million dimensions. Compared to the Web scale datasets with billions of documents in 2^{64} dimensions, *webspam* is relatively small and only uses 3-grams. Nevertheless, this example demonstrates that the set sizes (numbers of non-zeros), $f_i = |S_i|$, distribute in a wide range. Therefore, when we compare two sets, say S_1 and S_2 , we expect the ratio f_1/f_2 will often significantly deviate from 1.




Figure 1: Histograms of f_i (number of non-zeros in the i -th data vector) in *webspam* dataset.

Indeed, we computed the ratios f_1/f_2 for all pairs in *webspam*. Without loss of generality, we always assume $f_1 \geq f_2$. There are altogether 61 billion pairs with the mean $f_1/f_2 = 5.5$ and the standard deviation (std) = 9.5. Thus, we expect that $f_2/f_1 = 0.2 \sim 0.5$ is common and $f_2/f_1 < 0.1$ is also fairly frequent.

1.1 The 3-Cell Multinomial Problem

The standard estimator (2) is based on a binomial distribution. However, the problem really follows a 3-cell multinomial distribution. Define $z_1 = \min(\pi(S_1))$ and $z_2 = \min(\pi(S_2))$. The three probabilities are:

$$P_{=} = \Pr(z_1 = z_2) = \frac{a}{f_1 + f_2 - a} = R, \quad (4)$$

$$P_{<} = \Pr(z_1 < z_2) = \frac{f_1 - a}{f_1 + f_2 - a}, \quad (5)$$

$$P_{>} = \Pr(z_1 > z_2) = \frac{f_2 - a}{f_1 + f_2 - a} \quad (6)$$

These probabilities are easy to understand. For example, for the event $\{z_1 < z_2\}$, the size of sample space is $|S_1 \cup S_2| = f_1 + f_2 - a$ and the size of event space is $|S_1 - S_1 \cap S_2| = f_1 - a$, and hence $P_{<} = \frac{f_1 - a}{f_1 + f_2 - a}$.

We will show that the estimator solely based on $P_{=}$ (4) is optimal only when $f_1 = f_2$. Assuming $f_1 \geq f_2$, then (5) should not be used for the estimation task. The estimator based on $P_{>}$ (6) is superior to $P_{=}$ (6) when $f_1 \geq f_2 \approx a$. However, since we do not know a in advance, we must combine all three probabilities to ensure accurate estimates.

1.2 The Measure of Containment

The ratio $T = a/f_2$ (assuming $f_1 \geq f_2$) is known as the *containment*. It is possible that the resemblance R is small but the containment T is large. Note that $R = \frac{a}{f_1 + f_2 - a} \leq a/f_1 \leq f_2/f_1$. Thus, if, for example, $f_2/f_1 \leq 0.2$, then R has to be small, even when $a \approx f_2$ (which corresponds to $T \approx 1$).

While the literature on minwise hashing has mainly focused on the estimation of set resemblance, accurate estimation of set containment is also crucial to a number of different applications. For example, [9] uses both resemblance and containment estimates of the w -grams contained in text columns to characterize the similarity of database table contents in a tool that allows users to quickly understand database content. In a similar context, [31] tests the (estimated) level of containment between the distinct values contained in different (sets of) database columns to automatically detect foreign key constraints. [26] describes the use of (estimated) shingle containment in the context of cluster-based compression schemes. In the context of overlay networks, [6] uses the estimated containment (and resemblance) of the working sets of peers to coordinate between them, in turn reducing communication cost and complexity; because only small messages should be passed for coordination, this estimation has to be based on small synopses. The use of containment estimates in the context of peer-to-peer networking is discussed in [15].

1.3 b-Bit Minwise Hashing

The recent development of *b-bit minwise hashing* [22, 23, 20, 21] provides a solution to the (storage and computational) problem of *minwise hashing* by storing only the lowest b bits (instead of 64 bits) of each hashed value for a small b . [22] proved that using only $b = 1$ bit per hashed value can achieve at least a 21.3-fold improvement (in terms of storage) compared to using $b = 64$ bits if the target resemblance $R > 0.5$. This is a very encouraging result which may lead to substantial improvement in applications like (near)duplicate detection of Web pages [2].

On the other hand, when R is small, as shown in [22, 21], one might have to increase b in order to achieve an adequate accuracy without substantially increasing k , the number of permutations.

In fact, machine learning algorithms like SVM require (essentially) all pairwise similarities and it is expected that most pairs are not too similar. Our concurrent work [21] attempts to combine linear SVM [18, 28, 11, 16, 30] with b-bit hashing; and our initial experiments suggest that $b \geq 4$ (especially $b = 8$) is needed to achieve good performance.

In this paper, we will provide estimators for both the standard minwise hashing and b-bit minwise hashing.

2 Estimators for Minwise Hashing

Consider two sets $S_1, S_2 \in \Omega = \{0, 1, 2, \dots, D-1\}$. $f_1 = |S_1|$, $f_2 = |S_2|$. $a = |S_1 \cap S_2|$. We apply k random permutations $\pi_j : \Omega \rightarrow \Omega$, and record the minimums $z_{1,j} = \min(\pi_j(S_1))$, $z_{2,j} = \min(\pi_j(S_2))$, $j = 1$ to k . We will utilize the sizes of three disjoint sets:

$$k_{=} = |\{z_{1,j} = z_{2,j}, j = 1, 2, \dots, k\}| \quad (7)$$

$$k_{<} = |\{z_{1,j} < z_{2,j}, j = 1, 2, \dots, k\}| \quad (8)$$

$$k_{>} = |\{z_{1,j} > z_{2,j}, j = 1, 2, \dots, k\}| \quad (9)$$

Note that $E(k_{=}) = kP_{=}$, $E(k_{<}) = kP_{<}$, $E(k_{>}) = kP_{>}$, $Var(k_{=}) = kP_{=}(1-P_{=})$, etc. Thus, $\frac{k_{=}}{k}$, $\frac{k_{<}}{k}$, and $\frac{k_{>}}{k}$ are unbiased estimators of $P_{=}$ (4), $P_{<}$ (5), and $P_{>}$ (6), respectively. For the convenience of presentation, we estimate the intersection $a = (f_1 + f_2) \frac{R}{1+R}$:

$$\hat{a}_{=} = (f_1 + f_2) \frac{k_{=}/k}{1 + k_{=}/k} = \frac{(f_1 + f_2)k_{=}}{k + k_{=}} \quad (10)$$

$$\hat{a}_{<} = f_1 - f_2 \frac{k_{<}}{k - k_{<}} \quad (11)$$

$$\hat{a}_{>} = f_2 - f_1 \frac{k_{>}}{k - k_{>}}, \quad (12)$$

which are asymptotically (for large k) unbiased estimators of a . The variances are provided by Lemma 1.

Lemma 1

$$Var(\hat{a}_{=}) = \frac{1}{k} \frac{(f_1 + f_2 - a)^2 a (f_1 + f_2 - 2a)}{(f_1 + f_2)^2} + O\left(\frac{1}{k^2}\right) \quad (13)$$

$$Var(\hat{a}_{<}) = \frac{1}{k} \frac{(f_1 + f_2 - a)^2 (f_1 - a)}{f_2} + O\left(\frac{1}{k^2}\right) \quad (14)$$

$$Var(\hat{a}_{>}) = \frac{1}{k} \frac{(f_1 + f_2 - a)^2 (f_2 - a)}{f_1} + O\left(\frac{1}{k^2}\right) \quad (15)$$

Proof: The asymptotic variances can be computed by the “delta method” $Var(g(x)) \approx Var(x) [g'(E(x))]^2$ in a straightforward fashion. We skip the details. \square

2.1 The Maximum Likelihood Estimator

Lemma 1 suggests that the current standard estimator $\hat{a}_{=}$ may be severely less optimal when f_2/f_1 deviates

from 1. In fact, if we know $f_1 > f_2 \approx a$ (i.e., when the resemblance is small but the containment is large), we will obtain good results by using $\hat{a}_{>}$. The problem is that we do not know a in advance and hence we should resort to the maximum likelihood estimator (MLE).

Lemma 2 *The MLE, denoted by \hat{a}_{MLE} , is the solution to the following equation:*

$$k_{=} \frac{f_1 + f_2}{a} - k_{<} \frac{f_2}{f_1 - a} - k_{>} \frac{f_1}{f_2 - a} = 0 \quad (16)$$

which is asymptotically unbiased with the variance

$$Var(\hat{a}_{MLE}) = \frac{1}{k} \frac{(f_1 + f_2 - a)^2}{\frac{f_1 + f_2}{a} + \frac{f_2}{f_1 - a} + \frac{f_1}{f_2 - a}} + O\left(\frac{1}{k^2}\right) \quad (17)$$

Proof: The result follows from classical multinomial estimation theory. See Section 3.1. \square

2.2 Comparing MLE with Other Estimators

Figure 2 compares the ratios of the variances of estimators of a (only using the $O(\frac{1}{k})$ term of the variance). The top-left panel illustrates that when $f_2/f_1 < 0.5$ (which is common), the MLE \hat{a}_{MLE} can reduce the variance of the standard estimator $\hat{a}_{=}$ by a large factor. When the target containment $T = \frac{a}{f_2}$ approaches 1, the improvement can be as large as 100-fold.

The top-right panel of Figure 2 suggests that, if $f_2 \leq f_1$, then we should not use $\hat{a}_{<}$, because its variance can be magnitudes larger than the variance of the MLE. The bottom-left panel confirms that if we know the containment is very large (close to 1), then we will do well by using $\hat{a}_{>}$ which is simpler than the MLE. The problem is of course that we do not know a in advance and hence we may still have to use the MLE. The bottom-right panel verifies that $\hat{a}_{>}$ is significantly better $\hat{a}_{<}$.

2.3 Experiment

For the purpose of verifying the theoretical improvements, we use two pairs of sets corresponding to the occurrences of four common words (“A – TEST” and “THIS – PERSON”) in a chunk of real world Web crawl data. Each (word) set is a set of document (Web page) IDs which contained that word at least once. For “A – THE”, the resemblance = 0.0524 and containment = 0.9043. For “THIS – PERSON”, the resemblance = 0.0903 and containment = 0.8440.

Figure 3 presents the mean square errors (MSE) of the estimates using $\hat{a}_{=}$ and \hat{a}_{MLE} . The results verify our theoretical predictions:




Figure 2: Variance ratios (the lower the better). Except the bottom-right panel, we compare the other three estimators, $\hat{a}_=$, $\hat{a}_<$ and $\hat{a}_>$, with the MLE \hat{a}_{MLE} .

- For pairs of low resemblance and high containment, the MLE \hat{a}_{MLE} provides significantly better (in these two cases, about an order of magnitude better) results than the standard estimator $\hat{a}_=$.
- The MLE is asymptotically unbiased. The small bias at small k (which is common for MLE in general) vanishes as k increases.
- The theoretical variances match the simulations.




Figure 3: A simulation study using two pairs of real-world vectors (of low resemblance and high containment) to verify (i) \hat{a}_{MLE} is significantly better than $\hat{a}_=$; and (ii) the theoretical variances match the simulations and the bias of the MLE vanishes as k increases.

3 b-Bit Minwise Hashing

b-Bit minwise hashing [22] stores each hashed value, e.g., $z_1 = \min(\pi(S_1))$, $z_2 = \min(\pi(S_2))$, using the lowest b bits instead of 64 bits. In this section, we will show that because the original b-bit minwise hashing only used part of the available information, it can be substantially improved.

We first define:

$u_{1,b}$ = the number formed by the lowest b bits of z_1

$u_{2,b}$ = the number formed by the lowest b bits of z_2 .

[22] derived the probability formula $\Pr(u_{1,b} = u_{2,b})$ by assuming $D = |\Omega|$ is large (which is virtually always satisfied in practice). We will also need to derive

$$P_{b,(t,d)} = \Pr(u_{1,b} = t, u_{2,b} = d), \quad t, d = 0, 1, 2, \dots, 2^b - 1$$

We follow the convention in [22] by defining

$$r_1 = \frac{f_1}{D}, \quad r_2 = \frac{f_2}{D}, \quad s = \frac{a}{D} \quad (18)$$

Instead of estimating a , we equivalently estimate s in the context of b-bit hashing. Lemma 3 provides the probability formulas as the basic tool.

Lemma 3 Assume D is very large.

$$\begin{aligned} & \Pr(u_{1,b} = t, u_{2,b} = d, t < d) \quad (19) \\ &= P_{<} \frac{r_2 [1 - r_2]^{d-t-1} [1 - (r_1 + r_2 - s)]^t (r_1 + r_2 - s)}{1 - [1 - r_2]^{2^b} [1 - (r_1 + r_2 - s)]^{2^b}} \\ &+ P_{>} \frac{r_1 [1 - r_1]^{t+2^b-d-1} [1 - (r_1 + r_2 - s)]^d (r_1 + r_2 - s)}{1 - [1 - r_1]^{2^b} [1 - (r_1 + r_2 - s)]^{2^b}} \end{aligned}$$

$$\begin{aligned} & \Pr(u_{1,b} = t, u_{2,b} = d, t > d) \quad (20) \\ &= P_{>} \frac{r_1 [1 - r_1]^{t-d-1} [1 - (r_1 + r_2 - s)]^d (r_1 + r_2 - s)}{1 - [1 - r_1]^{2^b} [1 - (r_1 + r_2 - s)]^{2^b}} \\ &+ P_{<} \frac{r_2 [1 - r_2]^{d+2^b-t-1} [1 - (r_1 + r_2 - s)]^t (r_1 + r_2 - s)}{1 - [1 - r_2]^{2^b} [1 - (r_1 + r_2 - s)]^{2^b}} \end{aligned}$$

$$\begin{aligned} & \Pr(u_{1,b} = t, u_{2,b} = t) \quad (21) \\ &= \left(R + P_{<} \frac{r_2 [1 - r_2]^{2^b-1}}{1 - [1 - r_2]^{2^b}} + P_{>} \frac{r_1 [1 - r_1]^{2^b-1}}{1 - [1 - r_1]^{2^b}} \right) \times \\ &+ \frac{[1 - (r_1 + r_2 - s)]^t (r_1 + r_2 - s)}{1 - [1 - (r_1 + r_2 - s)]^{2^b}} \end{aligned}$$

Proof: See Appendix A. \square

Therefore, we encounter a multinomial probability estimation problem with each cell probability being a function of s . Note that the total number of cells, i.e., $2^b \times 2^b$, is large especially when b is not small.

In addition to $P_{b,(t,d)}$, we also define the following three

probability summaries analogous to $P_{=}$, $P_{<}$, and $P_{>}$.

$$\begin{aligned} P_{b,=} &= \Pr(u_{1,b} = u_{2,b}) = \sum_{t=0}^{2^b-1} P_{b,(t,t)} \\ P_{b,<} &= \Pr(u_{1,b} < u_{2,b}) = \sum_{t<d}^{2^b-1} P_{b,(t,d)} \\ P_{b,>} &= \Pr(u_{1,b} > u_{2,b}) = \sum_{t>d}^{2^b-1} P_{b,(t,d)} \end{aligned}$$

Suppose we conduct k permutations. We define the observed counts, $k_{b,(t,d)}$, $k_{b,=}$, $k_{b,<}$, and $k_{b,>}$, which correspond to $P_{b,(t,d)}$, $P_{b,=}$, $P_{b,<}$, and $P_{b,>}$, respectively. Note that $k = \sum_{t,d} k_{b,(t,d)}$.

[22] only used $P_{b,=}$ to estimate R (and hence also s). We expect to achieve substantial improvement if we can take advantage of the matrix of probabilities $P_{b,(t,d)}$. Here, we first review some basic statistical procedure for multinomial estimation and the classical (asymptotic) variance analysis.

3.1 Review Classical Multinomial Estimation

Consider a table with m cells, each of which is associated with a probability $q_i(\theta)$, $i = 1, 2, \dots, m$. Here we assume the probability q_i is parameterized by θ (for example, the s in our problem), and the task is to estimate θ . Suppose we draw k i.i.d. samples and the number of observations from the i -th cell is k_i , $\sum_{i=1}^m k_i = k$. The joint log-likelihood is proportional to

$$l(\theta) = \sum_{i=1}^m k_i \log q_i(\theta). \quad (22)$$

The maximum likelihood estimator (MLE), which is optimal or asymptotically (for large k) optimal in terms of the variance, is the solution $\hat{\theta}_{MLE}$ to the MLE equation $l'(\theta) = 0$, i.e.,

$$l'(\theta) = \sum_{i=1}^m k_i \frac{q'_i(\theta)}{q_i(\theta)} = 0, \quad (23)$$

solving which often requires a numerical procedure. For one-dimension problems as in our case, the numerical procedure is straightforward.

The estimation variance of $\hat{\theta}_{MLE}$ is related to the

Fisher Information $I(\theta) = -E(l''(\theta))$:

$$\begin{aligned} I(\theta) &= -E(l''(\theta)) = -\sum_{i=1}^m E(k_i) \frac{q''_i q_i - [q'_i]^2}{q_i^2} \\ &= -\sum_{i=1}^m k q_i \frac{q''_i q_i - [q'_i]^2}{q_i^2} = -k \sum_{i=1}^m \frac{q''_i q_i - [q'_i]^2}{q_i} \end{aligned} \quad (24)$$

$$\begin{aligned} \text{Var}(\hat{\theta}_{MLE}) &= \frac{1}{I(\theta)} + O\left(\frac{1}{k^2}\right) \\ &= \frac{1}{k} \frac{1}{\sum_{i=1}^m \frac{[q'_i]^2}{q_i} - q''_i} + O\left(\frac{1}{k^2}\right) \end{aligned} \quad (25)$$

For b -bit hashing, since we have $2^b \times 2^b$ cells with probabilities $P_{b,(t,d)}$, we can either use the full (entire) probability matrix or various reduced forms by grouping (collapsing) cells (e.g., $P_{b,=}$, $P_{b,<}$, and $P_{b,>}$) to ease the burden of numerically solving the MLE equation (23).

3.2 Five Levels of Estimators for s

We first introduce the notation for the following five estimators of s :

1. $\hat{s}_{b,f}$ denotes the full MLE solution by using all $m = 2^b \times 2^b$ cell probabilities $P_{b,(t,d)}$, $t, d = 0, 1, 2, \dots, 2^b - 1$. This estimator will be most accurate and computationally most intensive.
2. $\hat{s}_{b,do}$ denotes the MLE solution by using $m = 2^b + 2$ cells which include the 2^b diagonal probabilities $P_{b,(t,t)}$, $t = 0, 1, \dots, 2^b - 1$ and two summaries of the off-diagonals: $P_{b,<} = \sum_{t<d} P_{b,(t,t)}$ and $P_{b,>} = \sum_{t>d} P_{b,(t,t)}$.
3. $\hat{s}_{b,d}$ denotes the MLE solution by using $m = 2^b + 1$ cells which include the 2^b diagonal probabilities $P_{b,(t,t)}$, $t = 0, 1, \dots, 2^b - 1$ and the sum of the rest, i.e., $P_{b,<} + P_{b,>}$.
4. $\hat{s}_{b,3}$ denotes the MLE solution by using $m = 3$ cells which include the sum of the diagonals and two sums of the off-diagonals, i.e., $P_{b,=} = \sum_{t=0}^{2^b-1} P_{b,(t,t)}$, $P_{b,<}$, and $P_{b,>}$.
5. $\hat{s}_{b,=}$ denotes the MLE solution by using only $m = 2$ cells, i.e., $P_{b,=}$ and $1 - P_{b,=}$. This estimator requires no numerical solutions and is the one used in the original b -bit minwise hashing paper [22].

We compare the asymptotic variances of the other four estimators, $\hat{s}_{b,do}$, $\hat{s}_{b,d}$, $\hat{s}_{b,3}$, and $\hat{s}_{b,=}$, with the variance of the full MLE $\hat{s}_{b,f}$ in Figures 4 to 12. We consider $b = 8, 4, 6$, $r_1 = 0.8, 0.5, 0.2$ and the full ranges of $\frac{r_2}{r_1}$ and $\frac{s}{r_2}$ (which is the containment). Note that the

improvement of this paper compared to the previous standard practice (i.e., $\hat{s}_{b,=}$) is only reflected in the bottom-right panel of each figure. We present other estimators in the hope of finding one which is much simpler than the full MLE $\hat{s}_{b,f}$ and still retains much of the improvement. Our observations are:

- The full MLE $\hat{s}_{b,f}$, which uses a matrix of $2^b \times 2^b$ probabilities, can achieve substantial improvements (for example, 5- to 100-fold) compared to the standard practice $\hat{s}_{b,=}$, especially for cases of low resemblance and high containment.
- Two other estimators, $\hat{s}_{b,do}$ and $\hat{s}_{b,3}$ usually perform very well compared to the full MLE. $\hat{s}_{b,do}$ uses $2^b + 2$ cells and $\hat{s}_{b,3}$ uses merely 3 cells: the sum of the diagonals and the two sums of the off-diagonals. Therefore, we consider $\hat{s}_{b,3}$ is likely to be particularly useful in practice.

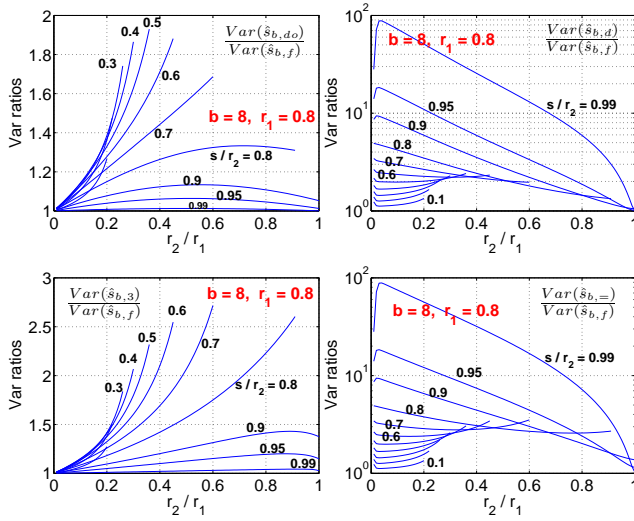


Figure 4: Given a contingency table of size $2^b \times 2^b$, we have defined five different estimators. We compare the theoretical variances of the other four estimators, $\hat{s}_{b,do}$, $\hat{s}_{b,d}$, $\hat{s}_{b,3}$, and $\hat{s}_{b,=}$, with the variance of the full MLE $\hat{s}_{b,f}$. The bottom-right panel measures the real improvement of this paper compared to the previous standard practice (i.e., $\hat{s}_{b,=}$). In this case, the variance ratios of about 10 to 100 are very substantial. The other three panels are for testing whether simpler estimators can still achieve substantial improvements. For example, both $\hat{s}_{b,do}$ and $\hat{s}_{b,3}$ (the left two panels) only magnify the variance of the full MLE by small factors compared to $\hat{s}_{b,=}$, and hence they might be good estimators in lieu of the quite sophisticated full MLE solution. In this figure, we consider $b = 8$ and $r_1 = 0.8$. Note that s/r_2 is the containment. The resemblance is upper bounded by r_2/r_1 .

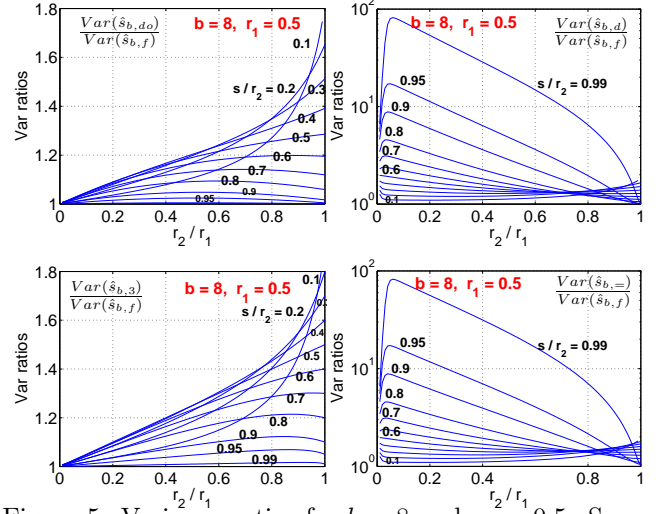


Figure 5: Variance ratios for $b = 8$ and $r_1 = 0.5$. See the caption of Figure 4 for more details.

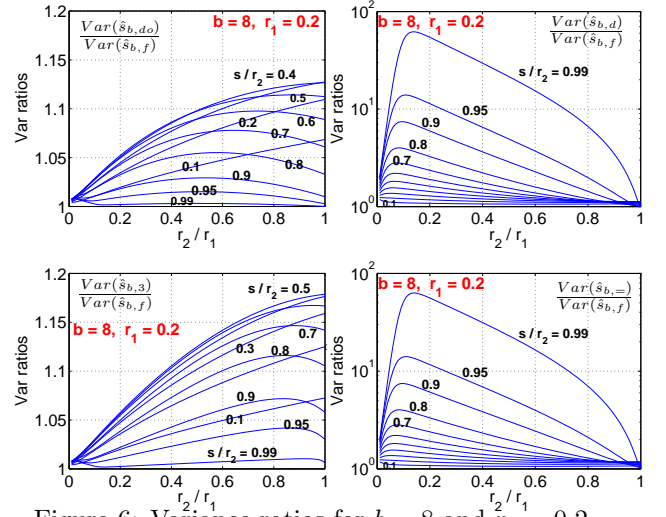


Figure 6: Variance ratios for $b = 8$ and $r_1 = 0.2$.

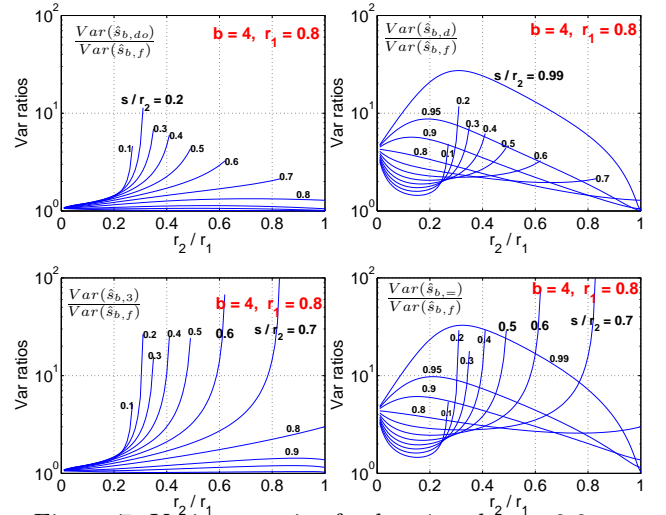


Figure 7: Variance ratios for $b = 4$ and $r_1 = 0.8$.

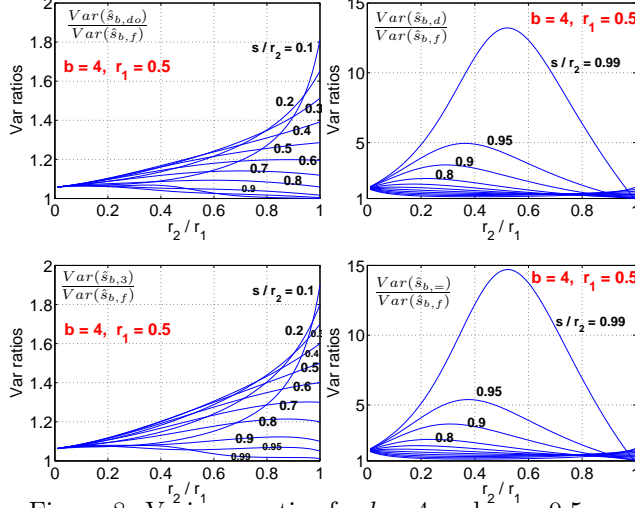


Figure 8: Variance ratios for $b = 4$ and $r_1 = 0.5$.

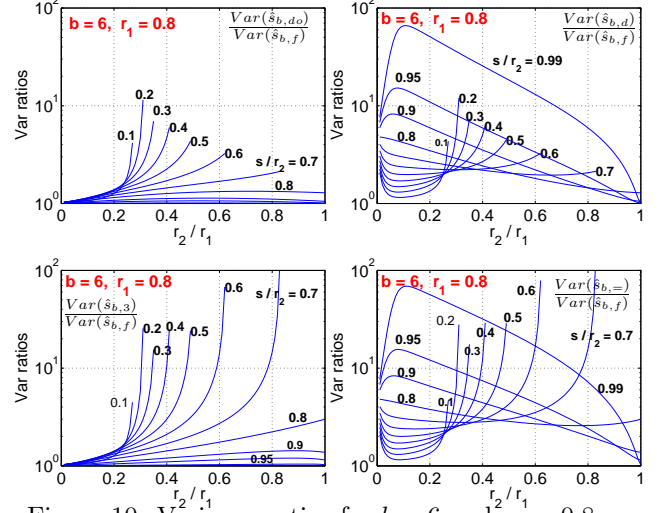


Figure 10: Variance ratios for $b = 6$ and $r_1 = 0.8$.




Figure 9: Variance ratios for $b = 4$ and $r_1 = 0.2$.

Our analysis has demonstrated that the much simpler estimator $\hat{s}_{b,3}$, which only uses 3 cells, is often remarkably accurate. We expect it will be used in practice. $\hat{s}_{b,3}$ involves three summary probabilities: $P_{b,=}$, $P_{b,<}$, and $P_{b,>}$. For efficient estimation, we will need to use more compact presentations instead of the double summation forms. Since we already know $P_{b,=}$ as derived in [22], we only need to derive $P_{b,<}$ and then $P_{b,>}$ follows by symmetry. After some algebra, we obtain

$$\begin{aligned}
 P_{b,<} &= \Pr(u_{1,b} < u_{2,b}) = \sum_{t < d} \Pr(u_{1,b} = t, u_{2,b} = d) \\
 &= \frac{1}{1 - [1 - r_2]^{2^b}} \frac{(r_1 - s)}{1 - [1 - (r_1 + r_2 - s)]^{2^b}} A \\
 &\quad + \frac{[1 - r_1]^{2^b - 1}}{1 - [1 - r_1]^{2^b}} \frac{(r_2 - s)}{1 - [1 - (r_1 + r_2 - s)]^{2^b}} B
 \end{aligned} \tag{26}$$

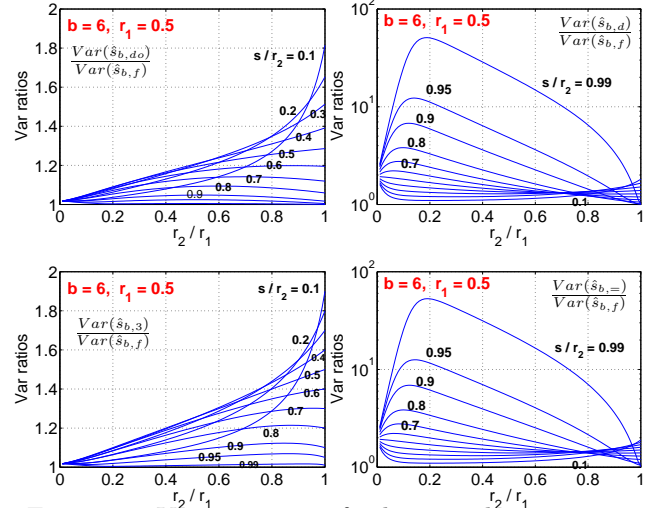


Figure 11: Variance ratios for $b = 6$ and $r_1 = 0.5$.

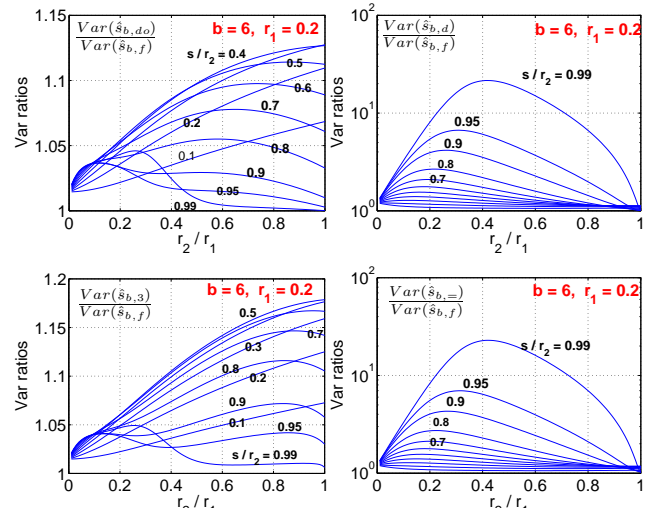


Figure 12: Variance ratios for $b = 6$ and $r_1 = 0.2$.

where

$$A = \frac{1 - [1 - (r_1 + r_2 - s)]^{2^b - 1}}{r_1 + r_2 - s} - \frac{[1 - r_2]^{2^b} - [1 - (r_1 + r_2 - s)]^{2^b - 1} [1 - r_2]}{r_1 - s}$$

$$B = \frac{[1 - r_1]^{2^b} - [1 - (r_1 + r_2 - s)]^{2^b}}{r_2 - s} [1 - r_1]^{1 - 2^b}$$

$$- \frac{[1 - (r_1 + r_2 - s)] - [1 - (r_1 + r_2 - s)]^{2^b}}{r_1 + r_2 - s}$$

4 Conclusion

Computing set or vector similarity is a routine task in numerous applications in machine learning, information retrieval, and databases. In Web scale applications, the method of minwise hashing is a standard technique for efficiently estimating similarities, by hashing each set (or equivalently binary vector) in a dictionary of size $|\Omega| = 2^{64}$ to about k hashed values ($k = 200$ to 500 is common). The standard industry practice is to store each hashed value using 64 bits. The recently developed b-bit minwise hashing stores only the lowest b bits with small b . b-Bit minwise hashing is successful in applications which care about pairs of high similarities (e.g., duplicate detection).

However, many applications involve computing all pairwise similarities (and most of the pairs are not similar). Furthermore, some applications really care about containment (e.g., the fraction that one object is contained by another) instead of resemblance. Interestingly, the current standard methods for minwise hashing and b-bit minwise hashing perform poorly for cases of low resemblance and high containment.

Our contributions in this paper include the statistically optimal estimator for standard minwise hashing and several new estimators for b-bit minwise hashing. For important scenarios (e.g., low resemblance and high containment), improvements of about an order of magnitude can be obtained. The full MLE solution for b-bit minwise hashing involves a contingency table of $2^b \times 2^b$ cells, which can be prohibitive if b is large. Our analysis suggests that if we only use 3 cells, i.e., the sum of the diagonals and the two sums of the off-diagonals, we can still achieve significant improvements compared to the current practice.

A Proof of Lemma 3

Consider two sets $S_1, S_2 \in \Omega = \{0, 1, 2, \dots, D - 1\}$. Apply a random permutation $\pi : \Omega \rightarrow \Omega$ on a S_1 , S_2 , and store the two minimums: $z_1 = \min(\pi(S_1))$, $z_2 = \min(\pi(S_2))$. Assuming $D \rightarrow \infty$, [22] provided two basic probability formulas:

$$\Pr(z_1 = i, z_2 = j, i < j) = r_2(r_1 - s) [1 - r_2]^{j-i-1} [1 - (r_1 + r_2 - s)]^i \quad (27)$$

$$\Pr(z_1 = i, z_2 = j, i > j) = r_1(r_2 - s) [1 - r_1]^{i-j-1} [1 - (r_1 + r_2 - s)]^j \quad (28)$$

We will also need to derive $\Pr(z_1 = i, z_2 = i)$. The exact expression is given by

$$\Pr(z_1 = i, z_2 = i) = \frac{\binom{D-i-1}{a-1} \binom{D-i-a}{f_1-a} \binom{D-i-f_1}{f_2-a}}{\binom{D}{a} \binom{D-a}{f_1-a} \binom{D-f_1}{f_2-a}}$$

$$= \frac{a(D-i-1)!(D-f_1-f_2+a)!}{D!(D-f_1-f_2+a-i)!}$$

$$= \frac{a}{D} \frac{\prod_{t=0}^{i-1} D - f_1 - f_2 + a - t}{\prod_{t=0}^{i-1} D - 1 - t}$$

For convenience, we introduce the following notation:

$$r_1 = \frac{f_1}{D}, \quad r_2 = \frac{f_2}{D}, \quad s = \frac{a}{D}.$$

Also, we assume D is large (which is virtually always satisfied in practice). We can obtain a reasonable approximation (analogous to the Poisson approximation of binomial):

$$\Pr(z_1 = i, z_2 = i) = s [1 - (r_1 + r_2 - s)]^i$$

To verify this, as expected,

$$\Pr(z_1 = z_2) = \sum_i \Pr(z_1 = i, z_2 = i)$$

$$= \sum_i s [1 - (r_1 + r_2 - s)]^i = \frac{s}{r_1 + r_2 - s} = R$$

Now we have tools to compute $\Pr(u_{1,b} = t, u_{2,b} = d)$, where $t, d \in \{0, 1, 2, 3, \dots, 2^b - 1\}$.

$$\Pr(u_{1,b} = t, u_{2,b} = d, t < d)$$

$$= \Pr(z_1 = t, t + 2^b, t + 2 \times 2^b, t + 3 \times 2^b, \dots,$$

$$z_2 = d, d + 2^b, d + 2 \times 2^b, d + 3 \times 2^b, \dots)$$

$$\Pr(z_1 = t, z_2 = d, d + 2^b, d + 2 \times 2^b, d + 3 \times 2^b, \dots)$$

$$= \sum_{j=0} \Pr(z_1 = t, z_2 = d + j \times 2^b, \dots)$$

$$= \sum_{j=0} r_2(r_1 - s) [1 - r_2]^{d+j \times 2^b - t - 1} [1 - (r_1 + r_2 - s)]^t$$

$$= r_2(r_1 - s) \frac{[1 - r_2]^{d-t-1}}{1 - [1 - r_2]^{2^b}} [1 - (r_1 + r_2 - s)]^t$$

$$\begin{aligned}
& \Pr(z_1 = t + 2^b, z_2 = d + 2^b, d + 2 \times 2^b, d + 3 \times 2^b, \dots) \\
&= \sum_{j=1} \Pr(z_1 = t + 2^b, z_2 = d + j \times 2^b, \dots) \\
&= \sum_{j=1} r_2(r_1 - s) [1 - r_2]^{d+j \times 2^b - t - 2^b - 1} [1 - (r_1 + r_2 - s)]^{t+2^b} \\
&= r_2(r_1 - s) \frac{[1 - r_2]^{d-t-1}}{1 - [1 - r_2]^{2^b}} [1 - (r_1 + r_2 - s)]^{t+2^b}
\end{aligned}$$

$$\begin{aligned}
& \Pr(u_{1,b} = t, u_{2,b} = d, t < d, z_1 < z_2) \\
&= \sum_{i=0} r_2(r_1 - s) \frac{[1 - r_2]^{d-t-1}}{1 - [1 - r_2]^{2^b}} [1 - (r_1 + r_2 - s)]^{t+i \times 2^b} \\
&= r_2(r_1 - s) \frac{[1 - r_2]^{d-t-1}}{1 - [1 - r_2]^{2^b}} \frac{[1 - (r_1 + r_2 - s)]^t}{1 - [1 - (r_1 + r_2 - s)]^{2^b}}
\end{aligned}$$

Next, we study

$$\Pr(u_{1,b} = t, u_{2,b} = d, t < d, z_1 > z_2)$$

$$\begin{aligned}
& \Pr(z_1 = d, \dots, z_2 = t + 2^b, t + 2 \times 2^b, t + 3 \times 2^b, \dots, t < d) \\
&= \sum_{j=1} r_2(r_1 - s) [1 - r_2]^{t+j \times 2^b - d - 1} [1 - (r_1 + r_2 - s)]^d \\
&= r_2(r_1 - s) \frac{[1 - r_2]^{t+2^b-d-1}}{1 - [1 - r_2]^{2^b}} [1 - (r_1 + r_2 - s)]^d
\end{aligned}$$

$$\begin{aligned}
& \Pr(z_1 = d + 2^b, \dots, z_2 = t + 2 \times 2^b, t + 3 \times 2^b, \dots, t < d) \\
&= \sum_{j=2} r_2(r_1 - s) [1 - r_2]^{t+j \times 2^b - d - 2^b - 1} [1 - (r_1 + r_2 - s)]^{d+2^b} \\
&= r_2(r_1 - s) \frac{[1 - r_2]^{t+2^b-d-1}}{1 - [1 - r_2]^{2^b}} [1 - (r_1 + r_2 - s)]^{d+2^b}
\end{aligned}$$

$$\begin{aligned}
& \Pr(z_1 = d, d + 2^b, d + 2 \times 2^b, \dots, \\
& \quad z_2 = t, t + 2^b, t + 2 \times 2^b, \dots, t < d, z_1 < z_2) \\
&= \sum_{i=0} r_2(r_1 - s) \frac{[1 - r_2]^{t+2^b-d-1}}{1 - [1 - r_2]^{2^b}} [1 - (r_1 + r_2 - s)]^{d+i \times 2^b} \\
&= r_2(r_1 - s) \frac{[1 - r_2]^{t+2^b-d-1}}{1 - [1 - r_2]^{2^b}} \frac{[1 - (r_1 + r_2 - s)]^d}{1 - [1 - (r_1 + r_2 - s)]^{2^b}}
\end{aligned}$$

By symmetry

$$\begin{aligned}
& \Pr(z_1 = t, t + 2^b, t + 2 \times 2^b, \dots, \\
& \quad z_2 = d, d + 2^b, d + 2 \times 2^b, \dots, t < d, z_1 > z_2) \\
&= r_1(r_2 - s) \frac{[1 - r_1]^{t+2^b-d-1}}{1 - [1 - r_1]^{2^b}} \frac{[1 - (r_1 + r_2 - s)]^d}{1 - [1 - (r_1 + r_2 - s)]^{2^b}}
\end{aligned}$$

Combining the results yields

$$\begin{aligned}
& \Pr(u_{1,b} = t, u_{2,b} = d, t < d) \\
&= r_2(r_1 - s) \frac{[1 - r_2]^{d-t-1}}{1 - [1 - r_2]^{2^b}} \frac{[1 - (r_1 + r_2 - s)]^t}{1 - [1 - (r_1 + r_2 - s)]^{2^b}} \\
&+ r_1(r_2 - s) \frac{[1 - r_1]^{t+2^b-d-1}}{1 - [1 - r_1]^{2^b}} \frac{[1 - (r_1 + r_2 - s)]^d}{1 - [1 - (r_1 + r_2 - s)]^{2^b}}
\end{aligned}$$

$$\begin{aligned}
& \Pr(u_{1,b} = t, u_{2,b} = d, t > d) \\
&= \Pr(u_{2,b} = d, u_{1,b} = t, d < t) \\
&= r_1(r_2 - s) \frac{[1 - r_1]^{t-d-1}}{1 - [1 - r_1]^{2^b}} \frac{[1 - (r_1 + r_2 - s)]^d}{1 - [1 - (r_1 + r_2 - s)]^{2^b}} \\
&+ r_2(r_1 - s) \frac{[1 - r_2]^{d+2^b-t-1}}{1 - [1 - r_2]^{2^b}} \frac{[1 - (r_1 + r_2 - s)]^t}{1 - [1 - (r_1 + r_2 - s)]^{2^b}}
\end{aligned}$$

Now, we need to compute $\Pr(u_{1,b} = t, u_{2,b} = t)$:

$$\begin{aligned}
& \Pr(u_{1,b} = t, u_{2,b} = t) \\
&= \sum_{i=0} \sum_{j=0} \Pr(z_1 = t + i \times 2^b, z_2 = t + j \times 2^b)
\end{aligned}$$

$$\begin{aligned}
& \Pr(u_{1,b} = t, u_{2,b} = t) \\
&= \Pr(u_{1,b} = t, u_{2,b} = t, z_1 = z_2) \\
&+ \Pr(u_{1,b} = t, u_{2,b} = t, z_1 \neq z_2)
\end{aligned}$$

$$\begin{aligned}
& \Pr(u_{1,b} = t, u_{2,b} = t, z_1 = z_2) \\
&= \sum_{i=0} \Pr(z_1 = z_2 = t + i \times 2^b, \dots) \\
&= \sum_{i=0} s [1 - (r_1 + r_2 - s)]^{t+i \times 2^b} \\
&= \frac{s [1 - (r_1 + r_2 - s)]^t}{1 - [1 - (r_1 + r_2 - s)]^{2^b}}
\end{aligned}$$

$$\begin{aligned}
& \Pr(u_{1,b} = t, u_{2,b} = t, z_1 < z_2) \\
&= \sum_{i=0} \sum_{j=i+1} \Pr(z_1 = t + i \times 2^b, z_2 = t + j \times 2^b) \\
&= \sum_{i=0} \sum_{j=i+1} r_2(r_1 - s) [1 - r_2]^{(j-i) \times 2^b - 1} [1 - (r_1 + r_2 - s)]^{t+i \times 2^b} \\
&= \sum_{i=0} r_2(r_1 - s) \frac{[1 - r_2]^{2^b-1}}{1 - [1 - r_2]^{2^b}} [1 - (r_1 + r_2 - s)]^{t+i \times 2^b} \\
&= r_2(r_1 - s) \frac{[1 - r_2]^{2^b-1}}{1 - [1 - r_2]^{2^b}} \frac{[1 - (r_1 + r_2 - s)]^t}{1 - [1 - (r_1 + r_2 - s)]^{2^b}}
\end{aligned}$$

$$\begin{aligned}
& \Pr(u_{1,b} = t, u_{2,b} = t, z_1 > z_2) \\
&= r_1(r_2 - s) \frac{[1 - r_1]^{2^b-1}}{1 - [1 - r_1]^{2^b}} \frac{[1 - (r_1 + r_2 - s)]^t}{1 - [1 - (r_1 + r_2 - s)]^{2^b}}
\end{aligned}$$

Combining the results yields

$$\begin{aligned}
& \Pr(u_{1,b} = t, u_{2,b} = t) \\
&= \frac{s [1 - (r_1 + r_2 - s)]^t}{1 - [1 - (r_1 + r_2 - s)]^{2^b}} \\
&+ r_2(r_1 - s) \frac{[1 - r_2]^{2^b-1}}{1 - [1 - r_2]^{2^b}} \frac{[1 - (r_1 + r_2 - s)]^t}{1 - [1 - (r_1 + r_2 - s)]^{2^b}} \\
&+ r_1(r_2 - s) \frac{[1 - r_1]^{2^b-1}}{1 - [1 - r_1]^{2^b}} \frac{[1 - (r_1 + r_2 - s)]^t}{1 - [1 - (r_1 + r_2 - s)]^{2^b}}
\end{aligned}$$

Finally, we re-write the probabilities in terms of $R = P_+$, $P_<$, and $P_>$ whenever possible:

$$\begin{aligned}
& \Pr(u_{1,b} = t, u_{2,b} = d, t < d) \\
&= P_< \frac{r_2 [1 - r_2]^{d-t-1}}{1 - [1 - r_2]^{2^b}} \frac{[1 - (r_1 + r_2 - s)]^t (r_1 + r_2 - s)}{1 - [1 - (r_1 + r_2 - s)]^{2^b}} \\
&+ P_> \frac{r_1 [1 - r_1]^{t+2^b-d-1}}{1 - [1 - r_1]^{2^b}} \frac{[1 - (r_1 + r_2 - s)]^d (r_1 + r_2 - s)}{1 - [1 - (r_1 + r_2 - s)]^{2^b}}
\end{aligned}$$

$$\begin{aligned}
& \Pr(u_{1,b} = t, u_{2,b} = d, t > d) \\
&= P_> \frac{r_1 [1 - r_1]^{t-d-1}}{1 - [1 - r_1]^{2^b}} \frac{[1 - (r_1 + r_2 - s)]^d (r_1 + r_2 - s)}{1 - [1 - (r_1 + r_2 - s)]^{2^b}} \\
&+ P_< \frac{r_2 [1 - r_2]^{d+2^b-t-1}}{1 - [1 - r_2]^{2^b}} \frac{[1 - (r_1 + r_2 - s)]^t (r_1 + r_2 - s)}{1 - [1 - (r_1 + r_2 - s)]^{2^b}}
\end{aligned}$$

$$\begin{aligned}
& \Pr(u_{1,b} = t, u_{2,b} = t) \\
&= \left(R + P_< \frac{r_2 [1 - r_2]^{2^b-1}}{1 - [1 - r_2]^{2^b}} + P_> \frac{r_1 [1 - r_1]^{2^b-1}}{1 - [1 - r_1]^{2^b}} \right) \times \\
&+ \frac{[1 - (r_1 + r_2 - s)]^t (r_1 + r_2 - s)}{1 - [1 - (r_1 + r_2 - s)]^{2^b}}
\end{aligned}$$

References

- [1] Michael Bendersky and W. Bruce Croft. Finding text reuse on the web. In *WSDM*, pages 262–271, Barcelona, Spain, 2009.
- [2] Andrei Z. Broder. On the resemblance and containment of documents. In *the Compression and Complexity of Sequences*, pages 21–29, Positano, Italy, 1997.
- [3] Andrei Z. Broder, Moses Charikar, Alan M. Frieze, and Michael Mitzenmacher. Min-wise independent permutations (extended abstract). In *STOC*, pages 327–336, Dallas, TX, 1998.
- [4] Andrei Z. Broder, Steven C. Glassman, Mark S. Manasse, and Geoffrey Zweig. Syntactic clustering of the web. In *WWW*, pages 1157 – 1166, Santa Clara, CA, 1997.
- [5] Gregory Buehrer and Kumar Chellapilla. A scalable pattern mining approach to web graph compression with communities. In *WSDM*, pages 95–106, Stanford, CA, 2008.
- [6] John Byers, Jeffrey Considine, Michael Mitzenmacher, and Stanislav Rost. Informed Content Delivery across Adaptive Overlay Networks. In *SIGCOMM*, pages 47–60, 2002.
- [7] Ludmila Cherkasova, Kave Eshghi, Charles B. Morrey III, Joseph Tucek, and Alistair C. Veitch. Applying syntactic similarity algorithms for enterprise information management. In *KDD*, pages 1087–1096, Paris, France, 2009.
- [8] Flavio Chierichetti, Ravi Kumar, Silvio Lattanzi, Michael Mitzenmacher, Alessandro Panconesi, and Prabhakar Raghavan. On compressing social networks. In *KDD*, pages 219–228, Paris, France, 2009.
- [9] T. Dasu, T. Johnson, S. Muthukrishnan, and V. Shkapenyuk. Mining database structure; or, how to build a data quality browser. In *SIGMOD*, pages 240–251, Madison, WI, 2002.
- [10] Yon Dourisboure, Filippo Geraci, and Marco Pellegrini. Extraction and classification of dense implicit communities in the web graph. *ACM Trans. Web*, 3(2):1–36, 2009.
- [11] Rong-En Fan, Kai-Wei Chang, Cho-Jui Hsieh, Xiang-Rui Wang, and Chih-Jen Lin. Liblinear: A library for large linear classification. *Journal of Machine Learning Research*, 9:1871–1874, 2008.
- [12] Dennis Fetterly, Mark Manasse, Marc Najork, and Janet L. Wiener. A large-scale study of the evolution of web pages. In *WWW*, pages 669–678, Budapest, Hungary, 2003.
- [13] George Forman, Kave Eshghi, and Jaap Suermondt. Efficient detection of large-scale redundancy in enterprise file systems. *SIGOPS Oper. Syst. Rev.*, 43(1):84–91, 2009.
- [14] Sreenivas Gollapudi and Aneesh Sharma. An axiomatic approach for result diversification. In *WWW*, pages 381–390, Madrid, Spain, 2009.
- [15] Systems Abhishek Gupta, Abhishek Gupta, Divyakant Agrawal, and Amr El Abbadi. Approximate Range Selection Queries in Peer-to-Peer. In *CIDR*, 2002.
- [16] Cho-Jui Hsieh, Kai-Wei Chang, Chih-Jen Lin, S. Sathya Keerthi, and S. Sundararajan. A dual coordinate descent method for large-scale linear svm. In *Proceedings of the 25th international conference on Machine learning*, ICML, pages 408–415, 2008.
- [17] Nitin Jindal and Bing Liu. Opinion spam and analysis. In *WSDM*, pages 219–230, Palo Alto, California, USA, 2008.
- [18] Thorsten Joachims. Training linear svms in linear time. In *KDD*, pages 217–226, Pittsburgh, PA, 2006.

- [19] Konstantinos Kalpakis and Shilang Tang. Collaborative data gathering in wireless sensor networks using measurement co-occurrence. *Computer Communications*, 31(10):1979–1992, 2008.
- [20] Ping Li and Arnd Christian König. Theory and applications b-bit minwise hashing. In *Commun. ACM*, 2011.
- [21] Ping Li, Anshumali Shrivastava, Joshua Moore, and Arnd Christian König. Hashing algorithms for large-scale learning. Technical report.
- [22] Ping Li and Arnd Christian König. b-bit minwise hashing. In *WWW*, pages 671–680, Raleigh, NC, 2010.
- [23] Ping Li, Arnd Christian König, and Wenhao Gui. b-bit minwise hashing for estimating three-way similarities. In *NIPS*, Vancouver, BC, 2010.
- [24] Gurmeet Singh Manku, Arvind Jain, and Anish Das Sarma. Detecting Near-Duplicates for Web-Crawling. In *WWW*, Banff, Alberta, Canada, 2007.
- [25] Marc Najork, Sreenivas Gollapudi, and Rina Panigrahy. Less is more: sampling the neighborhood graph makes salsa better and faster. In *WSDM*, pages 242–251, Barcelona, Spain, 2009.
- [26] Zan Ouyang, Nasir D. Memon, Torsten Suel, and Dimitre Trendafilov. Cluster-Based Delta Compression of a Collection of Files. In *WISE*, pages 257–268, 2002.
- [27] Sandeep Pandey, Andrei Broder, Flavio Chierichetti, Vanja Josifovski, Ravi Kumar, and Sergei Vassilvitskii. Nearest-neighbor caching for content-match applications. In *WWW*, pages 441–450, Madrid, Spain, 2009.
- [28] Shai Shalev-Shwartz, Yoram Singer, and Nathan Srebro. Pegasos: Primal estimated sub-gradient solver for svm. In *ICML*, pages 807–814, Corvallis, Oregon, 2007.
- [29] Tanguy Urvoy, Emmanuel Chauveau, Pascal Filoche, and Thomas Lavergne. Tracking web spam with html style similarities. *ACM Trans. Web*, 2(1):1–28, 2008.
- [30] Hsiang-Fu Yu, Cho-Jui Hsieh, Kai-Wei Chang, and Chih-Jen Lin. Large linear classification when data cannot fit in memory. In *KDD*, pages 833–842, 2010.
- [31] Meihui Zhang, Marios Hadjieleftheriou, Beng Chin Ooi, Cecilia M. Procopiuc, and Divesh Srivastava. On Multi-column Foreign Key Discovery. *Proc. VLDB Endow.*, 3:805–814, September 2010.